

DISTRIBUTED TRADING BUS ARCHITECTURE

BACKGROUND OF THE INVENTION

1. Field of the Invention.

[0001] The present invention relates, in general, to financial transaction software, and, more particularly, to software, systems and methods for processing transactions in a regulated exchange in a distributed fashion.

2. Relevant Background.

[0002] The present invention relates to systems, methods and a computer software architecture for trading in a futures trading marketplace. Futures and options trading are important techniques for coping with the price uncertainty of a free market. Price uncertainty creates risks and opportunities. Futures and options markets provide a forum for commercial interests in a commodity to hedge against price risk by transferring that risk to those more willing and able to bear it, or to those commercial interests with inverse risk profiles. An active futures market provides a readily available, widely accepted reference price for the underlying commodity, thereby improving the efficiency of the overall market. Futures can also be used for investment purposes to mitigate pure financial risks and/or pursue purely financial goals.

[0003] A futures contract is a contract between a buyer and seller, whereby the buyer is obligated to take future delivery and the seller is obligated to provide future delivery of a fixed amount of a commodity at a predetermined price at a specified location. The contracts are standardized so that the price of the contract has strong relationship to the value, and the expected fluctuations in value, of the underlying commodity. Futures contracts are traded until a set point in time before the contract-specified delivery date and in many, if not most, cases the positions are closed before physical delivery takes place. Futures contracts are traded exclusively on regulated exchanges. These auction

markets include, in many cases, "open outcry" markets as well as electronic trading. A futures market or exchange provides a mechanism where the futures contracts themselves can be bought and sold, much like a stock exchange provides a mechanism in which ownership in business entities can be bought and sold.

5 **[0004]** These exchanges implement systems that create and manage accounts for buyers and sellers, enable market participants to communicate transaction information, and execute transactions in a reliable fashion. It is desirable to provide these systems in a manner that minimizes transaction overhead. Like any market, a regulated exchange strives to bring many buyers and sellers together. Generally speaking, active markets are
10 considered more efficient. In other words, the larger the number of participants and the larger the transaction volumes, the more likelihood exists that the market will result in fair prices for the products being exchanged. In turn, confidence that prices are fair leads to greater participation. Hence, an exchange's systems and software must support large transaction volumes and scale over time to handle variable transaction volume and
15 numbers of participants.

[0005] Another attribute of a strong exchange is the ability to provide buyers and sellers with sufficient information about completed transactions so that they can better value their own transactions. Capturing, processing, and delivering transaction information in real time increases confidence of buyers and sellers. Accordingly, the exchanges strive
20 to provide systems and methods that reliably capture and record transactions, and execute those transactions efficiently and precisely.

[0006] An exchange also attempts to control transaction risks. While market participants accept risks associated with the underlying products they are purchasing (e.g., commodities futures and options), they desire to lessen risks associated with the
25 marketplace itself. For example, in any transaction a potential exists that a buyer will not have sufficient funds to complete a purchase. To protect against this "counterparty credit risk", exchanges require that market participants maintain a certain level of liquid assets on deposit called margin. As soon as anyone buys or sells a futures contract, they must deposit with their clearing member an amount of money that the exchange
30 determines is sufficient to cover any one-day price move. As long as that person or firm

holds on to the contract, the exchange must maintain minimum margin funds on deposit for that position, with the contract holder depositing additional funds whenever the market moves against him.

5 [0007] Contract values change continuously, however, and so exchanges implement processes to periodically assess the current value of customers' holdings and adjust margin requirements accordingly. "Clearing" refers to the processes of registration and settlement of a trade that includes provisions for margin requirement and performance guarantee. The "settlement price" is the price established by the exchange settlement committee at the close of each trading session and is the official price to be used by the clearinghouse in determining net gains or losses as well as margin requirements. In an active exchange the process of daily clearing involves millions of computations and account postings that must be performed in a very few minutes. Accordingly, the software and systems required to support this exchange activity are very complex.

15 [0008] Open outcry trading takes place on a physical trading floor where brokers exchange bids and offers for futures contracts. Executed trades are then recorded by hand or entered into an electronic recording system. The completed trades are later sent to an external or internal clearinghouse to process the trades and issue appropriate reports to the futures exchange and its members. Futures markets are also maintained on electronic trading systems. These electronic trading systems allow entry of a bid or offer for a particular futures contract. One can also enter orders for combinations of futures orders, e.g., spread or strip orders. A spread order buys one or more futures contracts and sells one or more futures contracts simultaneously at a single "differential" price. A strip order buys (or sells) two or more futures contracts simultaneously at a stated differential price from the previous settlement prices for each contract in the strip or at an average, i.e., "flat," price. The orders are time stamped by the trading system as they are entered. The system then matches a bid with an appropriate offer. Bids and offers are matched in the sequence in which they are received, hence, a buyer does not select a particular seller. The trading system then generates appropriate information for the clearinghouse.

[0009] Automated trading systems are complex software systems that assist exchanges in their efforts to implement the various process required by the exchange. These systems implement a variety of functions such as order entry, order validation, buy-sell order matching and publishing information in a form that is useful to transaction clearing entities. All of these functions are performed in a reliable, auditable fashion so that customers of the exchange can be assured of consistency. Automated futures trading systems were initially developed to support conventional trading by capturing data from trades executed by floor brokers. Such systems include processes for recording executed trades, rejecting unacceptable trades, and clearing acceptable trades. However, these systems are working with executed trades and so did not include processes for automating many of the tasks that were handled by the floor brokers.

[0010] Electronic trading systems that receive orders over a network, on the other hand, desirably include processes for executing trades. Early systems included processes for receiving orders over a network, but then handled those orders through a floor broker in a traditional manner. Another requirement for electronic trading systems is that they interface with the existing trading systems that support floor-brokered trades. Because of this need some systems have simply tacked on an electronic transaction interface to an existing centralized trading application. However, such approaches impose the deficiencies and limitations of the centralized trading application onto the electronic trading application.

[0011] Transactions must either be completed in a manner that satisfies both the buyer and seller, or must be discarded. A partially completed transaction is not allowed. A partially completed transaction is one that satisfies the buyer, but not the seller, or a transaction that satisfies both the buyer and seller partially, but not fully (e.g., when only a portion of both of these orders is filled). This factor, in combination with the highly specialized nature of trading systems, has led to centralized software implementation. These prior systems comprise complex, tightly coupled components that would execute on one or more mainframe computers.

[0012] However, complex, centralized software systems tend to be expensive to obtain, inflexible, and difficult to maintain. The tight coupling and interdependence between

functional components leads to unexpected results when components are upgraded or added. When the systems are based on mainframe computer environments, the hardware acquisition and maintenance is also expensive. As a result, these systems are slow to adapt to newer technologies and to support new business initiatives.

5 **[0013]** Often times centralized trading systems do not scale well. In the field of trading system software, it is not acceptable to have a system that fails when faced with large transaction volume. Accordingly, centralized systems are built from the beginning to support the largest expected transaction volume, even though this level of activity may be rarely if ever used. When the trading activity approaches the designed limit, the
10 system must be replaced or extensively reworked in order to support the larger transaction volumes. It would be desirable to implement trading system software and systems that scale gracefully to avoid the expense and disruption associated with system replacement.

15 **[0014]** Centralized trading systems, also called trading platforms, tend to offer limited and proprietary interfaces to external systems. Because a centralized system implements all or nearly all of the trading functions internally, there is little need for communication with external systems. Operators can be specially trained to handle various activities such as order entry and report generation, and so the systems tend to have limited, difficult to use interfaces to both humans and other machines. However, there is an
20 increasing desire to allow trading firms to use front-end applications that best fit their trading requirements. It is therefore desirable to be able to implement trading systems with improved ability to communicate with a wide variety of software and third party systems.

25 **[0015]** In many fields of software development, complex software is now being implemented in the form of distributed systems. These systems include functional components with a high degree of logical separation so that each component is largely insulated from the actions of other components. Distributed systems leverage the improved abilities to communicate between system components so that complex functions can be implemented as many small components. These smaller components
30 are easier to design and maintain, and provide much greater flexibility in adapting to a

variety of computer hardware/operating system environments. Until now, however, distributed system architectures have not been applied to trading system software.

SUMMARY OF THE INVENTION

5 [0016] Briefly stated, the present invention involves a distributed trading system architecture for handling a plurality of order requests, each order request comprising parameters under which a participant will buy and/or sell a futures contract. A validator component is coupled to a messaging bus and has a first interface for receiving order request and an interface generating a validated order message on the messaging bus related to validated orders, wherein the validator implements processes for validating the
10 order requests. A risk allocation value (RAV) component is coupled to the messaging bus and has an interface for receiving validated order messages from the validator, wherein the RAV component implements processes for evaluating risk associated with an order should that order be completed. A match engine is coupled to the messaging bus and has an interface for receiving validated order messages from the RAV
15 component, wherein the match engine implements processes for matching orders based on the order-specified criteria. A persist component is coupled to the messaging bus and has an interface for receiving messages related to orders and trades, wherein the persist component implements processes for persistently storing information related to orders and trades.

20 **BRIEF DESCRIPTION OF THE DRAWINGS**

[0017] Fig. 1 shows a distributed trading architecture in which the present invention is implemented;

[0018] Fig. 2 shows a distributed trading system environment of Fig. 1 in greater detail;

25 [0019] Fig. 3 shows an alternative view of the distributed trading system of Fig. 1 in greater detail;

[0020] Fig. 4 illustrates entities and selected messaging relationships between various entities in accordance with an embodiment of the present invention;

[0021] Fig. 5 illustrates the system of Fig. 4 with new order messaging exchanges highlighted; and

5 [0022] Fig. 6 illustrates the system of Fig. 4 with trade-related messaging exchanges highlighted.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

10 [0023] The present invention is illustrated and described in terms of a distributed computing system. The particular examples identify a specific set of components useful in a futures and options exchange, however, many of the components and inventive features are readily adapted to other electronic trading environments. The specific examples described herein may teach specific protocols and/or interfaces, although it should be understood that the invention is readily extended to other protocols and interfaces in a predictable fashion.

15 [0024] Regulated exchanges such as the New York Mercantile Exchange (NYMEX), assignee of the present invention, Chicago Board of Trade (CBOT), MidAmerica Exchange (MIDAM), Chicago Mercantile Exchange (CME), Kansas City Board of Trade (KCBT), and others implement a trading platform such as system 100 in Fig. 1 to conduct futures contract exchanges in an orderly manner. Trading platform 100 accepts
20 orders, matches and records orders; collects and maintains margins; allocates margins according to the positions of the clearing members; matches open short with open long positions for delivery; allocates delivery notices; and generates trading and delivery statistics. A clearinghouse 117 acts as a fiscal transfer agent, transferring money from the margin funds of traders who have incurred a loss in the futures market on any given
25 day to the margin funds of traders who have generated a gain. The exchange's clearing members accept responsibility for all trades cleared through them, and share secondary responsibility for the liquidity of the Exchange's clearing operation.

[0025] Trading platform 100 includes components that implement the various exchange functions necessary to support both electronic futures trading and trades executed on the trading floor 111. Conventional trading floor trades are managed and cleared by various processes within the common trading services 107. Electronic orders are conducted from an order submission point 101. In the particular example, order submission point 101 exists outside of the exchange and trading platform 101, however, there is no reason why an order submission point could not also supplement or replace trading floor activities.

[0026] Electronic orders are conducted by way of a communication network 103, such as the Internet. It is contemplated that order submission points may take on a wide variety of application-specific designs to suit the needs of particular brokerages, investors, investment plans and the like. Order submission point(s) 101 communicate with electronic execution services 105. Electronic execution services 105 implement, among other services, order validation logic to ensure that orders comply with certain standards imposed by the exchange. For this reason it is important that all communications from order submission point 101 enters platform 100 through a suitable order validation service.

[0027] Another important function of the trading platform 100 is to gather and provide market data to both market participants and external third parties. Market data services 109 gather market data 113 as transactions are executed. The market data 113 includes raw data, formatted data, as well as derivative (i.e., analyzed) data. Market data 113 may be provided by any available protocol and may be provided immediately and/or after a delay.

[0028] An important feature of the present invention is the use of enterprise messaging architecture (EMA) bus 115 to exchange messages between components of platform 100. In a specific example, EMA bus 115 is based on a TIBCO® SmartSocket® messaging bus, although equivalent alternatives include middleware such as Java Messaging Services, Arujuna Message Service, MQSeries messaging products available from IBM, and SonicMQ provided by Sonic Software as well as other messaging products. Messages are sent and received via a publish/subscribe method. Publish and

subscribe enables components 105, 107 and 109 share up-to-date information reliably. Unlike a request/reply-messaging scheme, the publish/subscribe method does not require a direct request for a message to be made before the message is published to EMA bus 115.

5 **[0029]** Using EMA bus 115, the present invention defines a set of messages between the various components. For each message, one component is a publisher and one or more other components are subscribers. A configuration table associated with each component identifies the messages to which that component subscribes. When a message is published, all subscribing components receive the message and can act on it
10 concurrently and independently. To change the subscriber-publisher relationships so that more or fewer components respond to particular message types, all that is required is to update the configuration tables of the affected components.

[0030] In this manner, EMA 115 implements a user defined and dynamically definable messaging protocol between the various components, and enables the components to
15 operate in an independent, distributed fashion. In an exemplary embodiment, messages on EMA bus 115 comprise self-describing messages such as extensible markup language (XML) documents. An XML schema is defined for each message type, for example, so that every component can validate messages as complying with the XML schema. Moreover, the variety and information content of each message can be readily
20 updated using an XML schema.

[0031] Fig. 2 shows a distributed trading system environment of Fig. 1 in greater detail, particularly with respect to the system components that are not connected to the EMA bus 115. The consumers and producers of information such as orders and market data are organized into clusters in particular embodiments. These components are essentially
25 the user interface components.

[0032] Trading floor component 111 comprises, among other functions, trade entry mechanisms 201 and price reporting mechanisms 203. Although trades executed on a trading floor are executed using traditional processes, they must be entered into the trading platform 100 to be handled by common trading services 107. The data entry will

indicate, for example, price, customer ID, information about the floor broker, time, commodity, delivery date and place of delivery, or other information required/desired by the exchange. Common trading services 107 may include processes for verifying that the trade meets certain exchange criteria and complies with exchange rules such as margin requirements, trading procedures, position sizes, and other criteria. Trading floor 111 may also include price reporting terminals 203 which may be implemented as network connected terminals coupled to trading services component 107 or as wireless devices such as personal digital assistants (PDAs). Price reporting terminals 203 can be used to enter price information related to executed trades and to view market data information on the trading floor 111.

[0033] Fig. 2 illustrates a variety of devices that can be used to implement order submission points 101. A FIX gateway 205 is an example of a third party standard system to which the present invention may readily be adapted to support. The Financial Information eXchange (FIX) protocol is a messaging standard developed specifically for the real-time electronic exchange of securities transactions. FIX is a public-domain specification owned and maintained by FIX Protocol, Ltd. FIX is a relatively easy-to-implement, open standard for communicating financial information in a secure but nonproprietary way between two parties. Increasingly, implementations are available from vendors in the form of toolkits or FIX-enabled software packages. FIX gateway 205 can be written specifically for an application or purchased in the form of a toolset from a number of vendors.

[0034] Electronic trading client 207 represents an application using either open standards or proprietary design that implements an order entry interface using client software such as a web browser. In a particular example, electronic trading client 207 is implemented with NYMEX ClearPort™ trading software that runs on a web browser such as Microsoft Internet Explorer. ClearPort is a trademark or registered trademark of the New York Mercantile Exchange, Inc. in the United States and other Countries. An electronic trading client 207 can be readily adapted to provide a customized, easy to maintain order entry interface that provides an application-specific set of functionality, in addition to supporting communication with electronic execution services 105. The

NYMEX ClearPort technology network is an extremely flexible system offering market participants wide latitude in meeting specific needs. NYMEX ClearPort is an open system, which allows trading firms to customize front-end applications to best fit their trading requirements.

5 **[0035]** Order submission point 101 also includes a market data application programming interface (API) 209 which functions to receive market data from market services component 109. This market data may include user-specific and/or application specific information on the status of trades, failure notifications, execution notifications and the like. The market data may also include general market information, news,
10 analysis and the like that can be used in planning future trades. In a particular example, market data is provided to market data API 209 using open standards protocols and formats such as Java™ and/or XML.

[0036] Cluster 215 represents devices for using market data from market data services component 109. For example, wallboard systems are frequently used to convey
15 information on trading floor 111. Electronic wallboards on each trading floor display virtually all needed price information on each of the most actively traded commodities, including the opening range, the previous day's closing price, the day's price limits, the high and low trade prices of the day, the three most recent trade prices and/or best bid or offer, and the net change
20 from the settlement price. In addition, the high, low, last, and net change in prices on other exchanges are displayed.

[0037] Providing market data to vendors 213 is a significant activity for an exchange. Market data from a futures exchange is frequently used as a basis for many transactions involving similar commodities outside of the
25 exchange. Vendors 213 purchase the market data and provide it in either raw or analyzed form to subscribers. The value of this market data to vendors 213 is largely dependent on its timeliness. Accordingly, the close integration of market data services 109 to the real time market data

available in trading platform 100 allows data to be provided to vendors 213 in an exceptionally timely and reliable fashion.

[0038] In a specific embodiment, the market data is provided to cluster 215 from market data services 109 using an ITC 2.1 protocol. ITC 2.1 is a standard message format that was developed by the Inter-Exchange Technical Committee. It is used by the major U.S. futures exchanges. The present invention supports the use of ITC 2.1 messages format for distributing market data from the market data server to devices that require that format, such as wallboards, data vendors, personal digital assistants (PDAs), and a variety of other data users. One feature of the present invention is an ability to collect, process, and display wallboard information in a timely fashion by placing market data services 109 on EMA 115.

[0039] Fig. 3 illustrates the embodiment of Fig. 1 with detailed emphasis on the components 105 and 107 that are coupled to EMA bus 115. Common trading services component 107 includes trade management system 301 that receives trade information from trading floor 111 and verifies that the trade information satisfies business rules of the exchange. Verified trades are made final and passed to an electronic settlement system (ESS) component 307. ESS component 307 comprises a system that takes trade survey data from outside cooperative traders and brokers and calculates a proposed settlement price. These proposed settlement prices are typically subjected to an approval process such as review by a research department, expert review, or other processes for scrutiny and approval that are implemented by a particular exchange. Proposed settlement prices are published to EMA bus 115 for dissemination to various systems including the price vendor community.

[0040] In the specific implementation of Fig. 3, ESS 307 receives data about completed trades from trade management system 301 and calculates a settlement price automatically as an automated substitute for the conventional role played by settlement price committee. Current practice

relies upon the decisions of settlement price committee member(s), whom are also trading members of the exchange, to determine what the settlement prices should be. In other words, traders themselves decide the settlement prices, which creates a conflict of interest situation. In contrast, the present invention contemplates that ESS 307 can suggest a settlement price and an exchange employee (e.g., an expert or research personnel) will adjust and/or approve the suggested settlement prices and publish this information. When ESS 307 is used in this latter capacity, it may continue to receive survey data from cooperative traders (e.g., data related to OTC-traded commodities) and calculate suggested settlement prices for research's consideration and approval. The particular algorithms used to compute settlement prices, the choice of what data is considered in the settlement price computation, and the form in which the settlement price information is presented are matters that can be selected to meet the needs of a particular application.

[0041] The settlement price is the final price established by the exchange settlement committee at the close of each trading session as the official price to be used by the clearinghouse in determining net gains or losses, margin requirements, and the next day's price limits. Final trades are also communicated to persist object 317 so that executed trades are safely stored in persistent storage of database 319.

[0042] Data related to settlement prices and trades are output to clearing processes 303. Clearing processes 303 are designed for high-volume, high-capacity clearing and settlement of exchange-based transactions. The system processes reported trades and tracks positions continuously in real time or near real time, providing users with instantaneous information on trades, positions and risk exposure. The open or outstanding contracts are communicated to ticker plant 305. Open interests are contracts for which an individual or entity is obligated to the exchange because that individual or entity has not yet made an offsetting sale or purchase, an actual contract delivery, or, in the case of options, exercised the option. Ticker plant 305 comprises a

collection software processes that consolidate the raw market data and distribute that data with low latency as well as high integrity and reliability. Ticker plant 305 distributes the market data that it collects and conveys that data to, for example, to market data services 109 and/or price reporting terminals 203.

5 **[0043]** Most of the component processes in common trading services 107 correspond directly to processes that have been used conventionally to handle trades from trading floor 111. However, by coupling common trading services 107 to the EMA bus 115, these otherwise conventional processes are now able to use new processes such as the persist component 317 and risk allocation value component 313. In this manner,
10 common trading services can be augmented by distributed services implemented by objects that are reachable through the EMA bus 115, greatly improving the expandability of common trading services 107.

[0044] Contract generator 309 comprises processes that enable spread and strip markets. A "strip" involves simultaneous purchase (or sale) of an equal number of futures
15 positions in consecutive months. A twelve-month strip, also known as a calendar strip, consists of an equal number of futures contracts for each of twelve consecutive contract months. A "spread" refers to a variety of transactions that involve simultaneous purchase and sale of futures contracts for different months, different commodities, different grades of the same commodity, or other combination in which the two sides of
20 the spread offset each other in some way. Hence, a contract order in the case of strips and spreads involves multiple buy/sell transactions, all of which must be completed for the contract order to be satisfied. Contract generator 309 receives data related to outright futures and options (i.e., contracts that do not involve spreads and strips) from the clearing processes 303 via EMA bus 115.

25 **[0045]** Contract generator 309 groups the outright futures contracts into "tradeable" spreads and strips and populates the contract tables in EMA database 310 with the strip and spread contracts. EMA database 310 contains firm, user, product, contract, orders, and trades data. This data is populated by direct user interfaces (firms, users, products), automated processes (contract populator and contract generator 309) as well as by trades

and orders being written to the database due to trades and orders occurring in the system. The trading system loads the outright future, strip and spread contracts for trading.

5 [0046] The electronic execution services 105 shown on the right hand side of Fig. 3 include a number of functional components that illustrate the power of the distributed program architecture in accordance with the present invention. Each of the components within electronic execution services 105 can be implemented by, for example, a program object such as a Java™ object. Java is a trademark or registered trademark of Sun Microsystems Inc. in the United States and other countries. Many instances of the various objects can exist at any given point in time thereby enabling a high level of
10 concurrency in program execution. Each instance of a component is associated with a particular contract or cluster of contracts. In general, failure of one object will not impact other objects, although care is taken to ensure that all objects have a synchronized view of the state of each transaction that is being handled by the system.

15 [0047] The present invention may be implemented in a manner that supports multiple "contract clusters" or in a manner in which contract clustering is not involved. A contract cluster is essentially a group of contracts that are sufficiently related that offers to sell and bids to buy the contracts within a contract cluster can be matched against each other. For example, a "contract cluster" may comprise two contracts sharing a spread or strip relationship that requires that implied markets be generated between the
20 two markets. In some implementations it may improve performance and efficiency to create and maintain components such as match engine 315 and persist component 317 on a contract cluster basis rather than a single instance for all contracts. When contract clustering is used, a unique identification called the cluster ID is associated with each cluster. In the particular example, each message placed in the EMA bus is tagged with
25 the cluster ID. Each component is configured to recognize a particular cluster ID and ignore messages associated with other cluster IDs. Accordingly, when a new cluster is added or removed, the component configuration files are updated accordingly.

30 [0048] Validator 311 is the entry point for orders into the system, either from client 207, FIX interface 205, or an equivalent order entry front-end. Validator 311 performs various validations (exchange active, contract active, markets open, user assigned to

account, high/low limits etc.) and either passes on successful orders to risk allocation value (RAV) component 313 or 315. Errors are reported back to the client. Any number of validators 311 may be instantiated at any given time. In a typical embodiment, each client connection initiates the creation of a validator 311.

5 **[0049]** Validator 311 is a publisher for an "AddOrder" message that is subscribed to by risk allocation value component 313. An AddOrder message is placed on EMA bus 115 when an order request has been validated. An AddOrder message comprises an XML document that is validated against an AddOrder XML schema. The validator 311 subscribes to market state messages enabling the use of market state information when
10 validating orders. Validator 311 also subscribes to RejectOrder messages from RAV component 313 and/or match engine 315. Validator 311 includes processes for notifying the customer that submitted a rejected order of the rejection and optionally other data that describes reasons for the rejection.

15 **[0050]** Because the messages on EMA bus 115 are self-describing, it is not necessary for validator 311 to maintain state information for each transaction that passes through it. If an order is rejected, the associated RejectOrder message on the EMA bus 115 will include sufficient information to enable validator 311 to fulfill its function of generating appropriate notifications.

20 **[0051]** Risk allocation value component (RAV) component 313 includes processes that evaluate the affect of proposed transaction against clearing member rules such as margin requirements, position size, and the like. For example, a clearing member may establish an account and set an overall "RAV limit" in dollars terms, e.g., \$100,000, to control the total size of positions that the account may establish across all commodity contracts that the clearing member has permissioned the account to trade. The RAV limit works in a
25 similar fashion as does "original margin." In a specific implementation, upon entry of a new order, RAV component 313 accepts or rejects the order by considering the existing positions for the account identified on the order and any matchable order for that account. After determining a "worst case" trading scenario (i.e., a scenario that leads to the greatest RAV requirement), RAV component 313 calculates the RAV value
30 necessary to support the worst case position generated by the order request and, based on

the account's preset RAV limit of \$100,000, either accepts or rejects the order. In addition to the clearing member setting an overall RAV limit value and permissioning the account to trade in particular commodity contracts, a clearing member may set net long and/or net short position limits per commodity for the account. When either the RAV limit or the position limit values would be exceeded due to the entry of a new order, then the order will be rejected and it will not be matchable in the system. RAV component 313 is a subscriber to AddOrder messages from validator 311 as well as RejectOrder messages from the match engine 315. Upon receiving a RejectOrder message, RAV component 313 may take remedial action to note failure of the order. Such action may include allowing another order to pass to its match engine 315 as the risk allocation value computations can now ignore the failed order request.

[0052] Match engine 315 will reject orders for a variety of reasons such as duplicate order identifications, format problems, and the like. Assuming the nominal case of a well-formed order, match engine 315 passes all new orders, regardless of whether they are matched or not, to the persist component 317. This is performed using a PERSIST message on the EMA bus 115 for which match engine 315 is a publisher and persist component 317 is a subscriber. The Persist message includes information indicating whether the order was matched to another complementary order, so that the persist component 317 can be synchronized with the trade information, if necessary.

[0053] Match engine 315 should subscribe to EMA bus messages pertaining to high/low limits. Match engine 315 will need to round prices for implied transaction to high/low limits. When the limits change, match engine 315 recalculates the markets. When an implied bid is generated above the high limit, then match engine 315 rounds the resultant high limit bid price to the limit price for display purposes on trading screens and in market data. When the high limit price was increased via a price limit adjustment, then the match engine recalculates the high bid price that should be displayed to the marketplace. If the new high limit price is still less than the system-generated implied bid price, then the new high bid price would be the new limit price. If the implied bid price were less than the new high limit price, then the new high bid price is the system-generated implied bid price. Match engine 315 also publishes EMA bus

messages that convey market data such as best bid and best asked for the contract cluster. These market data messages are subscribed to by the ticker plant 305 and market data services component 109, for example.

5 [0054] In a particular embodiment, match engine 315 and persist component 317 are scaled at a contract cluster level of granularity. That is to say, an instance of match engine 315 and an instance of a persist component 317 are created for every contract cluster rather than for every contract. As noted hereinbefore, a "contract cluster" may comprise two contracts sharing a spread or strip relationship that requires that implied markets be generated between the two markets. When this happens, matching engine 10 315 must consider an order entered into one contract as potentially causing a match in the second contract, thereby requiring that the matching engine 315 consider the two contracts simultaneously when determining matches. When contract clustering is used, a contract can only belong in one cluster and many contracts can be in the same cluster.

15 [0055] For example, a single contract cluster might exist for contracts such as light sweet crude oil (CL), heating oil (HO), unleaded gas (HU), Brent Crude Oil (SC), Light Louisiana Sweet Crude Oil (LS), Mars blend crude oil (MB), West Texas Sour Crude Oil (TS), because there are CLHO, CLHU, HOHU, SCCL, TSCL, MBCL, LSCL spread markets that require implied markets to be generated. The four-letter market identifiers indicate a spread contract comprising two markets, each identified by a two-letter 20 identifier. Hence, the contract cluster includes all of the constituent outright futures contracts that are necessary to make up a spread transaction.

25 [0056] In a futures contract exchange all contracts for delivery of a specific product at a specific delivery point and delivery date will be together in a contract cluster. Tens, hundreds or thousands of contract clusters will exist at any given time. Contract clusters are associated with a unique identification, and messages on the EMA bus 115 include the relevant contract cluster ID. In this manner an order is readily directed to the match engine that is relevant to that order. A match engine instance will remain alive in the system for as long as the contract cluster exists. As each new contract cluster is added, a corresponding new match engine 315 is created. As a contract cluster rolls off (e.g., 30 when the underlying futures contract delivery date passes) the corresponding match

engine 315 is destroyed. When a contract cluster is added or rolled off, the configuration files for each of the system components are updated to reflect a current list of active contract clusters.

5 [0057] Although contract cluster can be made larger, searching for matches becomes incrementally more complex and time-consuming as match engine 315 must attempt to match more parameters. In the implementation described above, matching is a much simpler algorithmic process. Matching is performed on a first in first-match basis. Accordingly, given two bids at the same price and only one offer, the bid that was first in time to reach match engine 315 will be matched. The competing bid will go
10 unmatched until a subsequent satisfactory offer arrives at match engine 315.

[0058] Each order is preferably associated with a time in force value indicating whether the order is a day order, good until cancelled (GTC), or good until the end of a specific date (GTD). The match engine 315 may identify a matching order at any time during its existence for orders that remain active. When a match is found the match engine
15 publishes a TRADES EMA bus message to the persist component 317. Match engine 315 generates the trade sequence IDs and trade timestamp. In practice, a match may require two messages to persist component 317, one for the new order that caused the trade and one for the TRADES message. The persist component is responsible for synchronizing the persistence of orders and trades that change due to the same event. In
20 a particular embodiment, match engine 315 tags data with an event ID and orders tagged with the same event ID are matched, which enables the persist component 317 to know what records to synchronize and when.

[0059] Upon expiry of the time in force, the Persist process 317 will expire orders upon the close of a market, e.g., or at the beginning of a new trading day. The matching
25 engine 315 receives a message to cancel the expired order(s) and processes the request. At the point that the order is canceled by the matching engine due to it expiring, the order is removed from the memory allocated to match engine 315 handling the order.

[0060] Orders may define a spread. A spread is the simultaneous purchase and sale of futures contracts for different months, different commodities, or different grades of the

same commodity. A spread trade involves a "leg" component contract for each transaction. For example, one leg corresponds to the purchase and one leg corresponds to the sale. The present invention contemplates a leg pricing component 321 that prices "legs" (i.e., component contracts) involved in spread orders. Leg pricing component 321 may be implemented as a separate component or within match engine 315. When two spread orders match at a stated differential price between the two contracts involved in the spread, each component contract needs an "absolute" price so that it may be cleared independently of the other component contract to the spread trade. In the particular embodiment, leg pricing component 321 is stateless. Leg pricing component 321 includes processes to read settlement, last price and high/low limits from the database when needed. Leg pricing processes 321 receive data related to high/low limits in order to determine if trade is price discovery or not. Any number of leg pricing components 321 may be created to meet the needs of a particular application.

[0061] Market data server component 109 operates to publish market data to any interested (i.e., subscribing) party, including clients 101, price reporting service terminals 203, trade management system 301, and others. Market data server component 109 preferably filter trades before broadcasting market data (price discovery, in and out of price limits, and the like). In the examples, an instance of a market data server component 109 corresponds to a specific market data service. For example, "level 1 market data" includes best bid and best offer price information whereas "level 2 market data" includes all bid and offer price information regardless of whether the bid or offer prices are the best or not. Any number of market data server components 109 may exist. These components will typically exists for so long as the level of data that they manage is being used by an external consumer.

[0062] One advantage of a distributed processing architecture as described above is that failure or malfunction of one instance of one component is unlikely to affect other instances of that same component or instances of different components. In this manner, a badly behaved instance is isolated. In a system in which many thousands of transactions may be pending at any instant, the ability to shut down one or only a few component instances and restart them is a powerful feature.

5 [0063] A LOCKDOWN message published on EMA 115 by any component will causes the exchange to halt. This means that no order entry or order maintenance is allowed. Validator instances stop accepting new orders and modifications to orders. Client market data displays, if synchronized with orders/trades, will be left alone. In the event that the client market data display is not synchronized, the display is preferably cleared to avoid unintended reliance on inaccurate data. End users and operators are informed of the halt.

10 [0064] A manual 'Resynch' process (button on admin or restart application) may be used to respond to and/or recover from a LOCKDOWN due to a critical component failure. A resynch will restart/reload market data system 109, restart/reload M.E. 315, restart/reload RAV (if RAV had the issue) 313, and restart/reload O.V. 311. At the appropriate time after appropriate notice to traders, the Exchange can be un-halted ("reopened") and everyone can resume trading activity.

15 [0065] Desirably, a robust and reliable monitoring system is in place for these components, one that can detect real failures in components, not raise false alarms, and can perform automated 'tests' to verify the loss of one of these components (hardware dead, network dead, and the like). When one of these components fails, action must be taken to recover the system such that all the components in the system are in a consistent state.

20 [0066] When persist component 317 detects database 319 is not functioning, the persist component 317 publishes a LOCKDOWN message on EMA bus 115, triggering the LOCKDOWN processing outlined above. Operations and support personnel address the issue as soon as possible to ensure functionality of database 319. Appropriate safeguards in the design of database 319 including RAID-type data protections and
25 virtualized storage can make the likelihood of failure of database 319 extremely remote. Once the underlying issue is resolved a resynch is triggered, as outlined above.

[0067] Upon failure of a market data server 109 operations personnel are alerted. Operations personnel may manually monitor and/or use monitor processes to run tests to confirm the market data system 109 process is dead. When market data system 109 is

verified dead, a LOCKDOWN message can be automatically published by the monitor if it ran the tests, or by the operations personnel manually if they run the tests. A Resynch is triggered, as outlined above.

5 [0068] When match engine 315 goes down a LOCKDOWN message should be published to EMA bus 315. Operations personnel and support personnel solve the issue ASAP, which ultimately leads to having the match engine 315 back up and running. A resynch is triggered, as outlined above.

10 [0069] If an issue is detected on either RAV 313 running as a single instance, or on one or more of the RAV engines 313 running on many instances, and that are important enough to cause a halt in trading, the following steps should be performed. All markets should be halted to allow downward modifications to order volume only by traders/FIX clients. Operations personnel and support personnel solve the issue as soon as possible. Some or all of the RAV instances may be recovered to return the system to a running condition. Minor inaccuracies in RAV instances 313 can be tolerated in some
15 applications. At the appropriate time, after notice to traders, all markets may be reopened.

[0070] Upon failure of a validator 311, the exchange is halted to prevent all trading activity. Of course, if validator instances 311 cannot be created, no new orders can enter the system. Operations personnel and support personnel must solve the issue. At the
20 appropriate time after appropriate notice to traders, the exchange can be un-halted and trading activity can resume.

[0071] In competitive markets such as a futures exchange, it is important to implement procedure and rules to prevent publication of market data that reflects trades that in fact did not occur due to a customer cancellation or change, failure of a component, or other
25 anomaly. The architecture of the present invention makes market data available so quickly that this can become a significant issue. Accordingly, EMA bus 115 components preferably follow a "persist before publish" rule in specific circumstances. Each subscriber component will need the ability to send an error back to publisher component, with possible rollback processing as a result.

[0072] Whether the publishing of market data is to be synchronized with the successful persistence and publishing of order/trade data that caused the market data is optional. Market data synchronization effectively means the system does not publish any level 1 or level 2 market data that is caused by an order or trade that failed to be persisted and thus will not exist when the system components resynchronize off the databases subsequent to a LOCKDOWN, as described above. However, the market data synchronization of the successful persistence of orders and trades with the distribution of level 1 and 2 data is at the possible expense of throughput and performance.

[0073] To enable the system to keep a full audit trail of order status changes, for both the end user and compliance purposes, the match engine 315 reports at least a minimum amount of information required for the persist component 317 to create the audit trail. Additionally, match engine 315 should support "pass through data", i.e., information that is associated with an order (user, account, etc.) that is of no interest to the match engine 315, but is of definite interest to the persisting process at the end of the stream.

[0074] Optionally, validator component 311 could keep a record of incoming order changes requested by a client, and strike these from its list once it gets confirmation from EMA bus 115 that this order change persisted. This gives us the ability, should a LOCKDOWN message go out and the validator 311 is up and running, to report back to a client the specifics of any orders that did not make it into the system rather than a general message.

[0075] Fig. 4 illustrates a logical view of a system in accordance with the present invention with emphasis on the interconnectivity provided by EMA bus 115. Fig. 5 and Fig. 6 are similar to Fig. 4 but differ in that communications between components through EMA bus 115 are illustrated by bold lines. In Fig. 4 through Fig. 6, each of the arrowed lines indicates an EMA bus message. Each message may have any number of subscribers. In Fig. 4 through Fig. 6, an arrow pointing away from a component indicates that the component is a publisher of that message. Conversely, an arrow pointing towards a component indicates that the component is a subscriber to that message.

[0076] Any number and variety of messages may be defined for a particular trading environment, and only a subset of messages are illustrated herein. Moreover, each message is preferably self-describing such as may be implemented using XML documents as the message carrier. This enables new message types to be defined quickly and easily without impacting existing message types nor requiring changes to components. Typically, each subscriber takes some action in response to hearing a message to which it subscribes.

[0077] Table 1 below includes an exemplary subset of EMA bus messages used in an implementation of the present invention. In table 1, the column labeled "subject" includes a name given to a particular message type, the column labeled "publisher" identifies the component that generates the bus message, "subscriber" indicates the components that receive or subscribe to that message type, and the "comments" column provides a brief description of the purpose of the message types. Messages including a <clusterid> designation indicate messages that, when contract clustering is used, will be specifically addressed to subscriber components associated with a particular cluster ID. In these cases, a unique instance of the subscribing component exists for each unique contract cluster ID. When contract clustering is not used, the "clusterid" value is not required in the message.

Subject	Pub	Sub	Comment
/cpt/orderrequest/<clustid>	V	RAV	A new order from the Validator to RAV
/cpt/orderaccept/<clustid>	RAV	ME	RAV accepts, sends to ME
/cpt/orderreject	RAV	V	RAV rejects, sends to V.
/cpt/orderactive/<clustid>	ME	Persist	M.E. accepted new order, sends to Persist
/cpt/error/order	ME	V RAV Ops MDS.	M.E. rejects order due to error (dupe order ID, etc.) RAV listens to undo order, V listens to send error to client
/cpt/marketdata/level2/<clustid>	ME	T.P.	
/cpt/mds/level2/<clustid>	T.P.	MDS	
/cpt/ordernewpersist/<clustid>	Persist	MDS	Persist has stored the order, sends to MDS
/cpt/marketdata/level1/<clustid>	T.P.	MDS	T.P. creates Level 1 data, sends to MDS
/cpt/mds/orderupdate	Persist	MDS	
/cpt/tradepersist	Persist	RAV MDS Trans	
/cpt/reallocationsrequest	V	RAV	
/cpt/reallocationsaccept	RAV	Persist	
/cpt/reallocationsreject	RAV	V	
/cpt/reallocationspersist	Persist	Trans MDS	
/cpt/trades	ME	Persist	
/cpt/setcontractstates	S.G.	Persist	Sending state group updates to be persisted
/cpt/me/setcontractstates	Persist	ME	After persisting state change, sent to ME
/cpt/tp/setcontractstates	ME	T.P. RAV	State group updates after ME has done with them
/cpt/mds/setcontractstates	T.P.	MDS	
/cpt/setsystemstates/eod	S.G.	RAV	
/cpt/setsystemstates/sod	S.G.	T.P.	
/cpt/orderinactivenew	V	Persist	New (inactive) order, straight to persist
/cpt/ordermodify	V	RAV	Modified order
/cpt/ordermodifyrequest	RAV	ME	
/cpt/ordermodified	ME	Persist	
/cpt/orderinactivemodify	V	Persist	Modified (inactive) order, straight to persist
/cpt/orderspull	V	ME	Cancel / deactivate
/cpt/orderspulled	ME	Persist RAV	

/cpt/ordersinactivepull	V	Persist	Cancel an inactive order, straight to persist
/cpt/ordermodifyreject/rav	RAV	V	
/cpt/ordermodifyreject/me	ME	V RAV	
/cpt/orderspullreject	ME	V	
/cpt/mds/positionupdate	Persist	MDS	
/cpt/voidtrade	V	Persist	
/cpt/voidpersist	Persist	RAV MDS Trans	
/cpt/tp/voidpricereport	Persist	T.P.	
/settlements	TMS ESS	ME Persist TP	
/cpt/lockdown	Monitor	V	
/cpt/unlock	Monitor	V	

TABLE 1

Key: V=Validator Component

RAV=RAV Component

M.E.=Match Engine Component

T.P.=Ticker Plant Component

Persist=Persist Component

S.G.= State Generator

MDS=Market Data Services Component

T.M.S.=Trade Management System

5
10
15
[0078] Referring to Fig. 4, validator 311 publishes a message (e.g., an "AddOrder" message described above) upon validating a received order request. Validator 311 also publishes messages related to order changes such as order volume changes and order price changes in response to receiving validated order change requests. Validator 311 also publishes a message in response to a new trigger. Although not shown in Fig. 4, validator 311 also subscribes to various messages relating to order rejections by components such as RAV 313 and match engine 315.

20
[0079] RAV component 313 subscribes to messages relating to new orders, order changes, and executed trades. RAV component 313 publishes messages related to new orders that satisfy the risk allocation valuation rules implemented by RAV component 313. With respect to order changes and executed trades, RAV component 313 uses the information to adjust the information it maintains about each account, such as position

sizes and margin information, and the like. In the case of an executed trade, for example, RAV 313 does not have a defined role in accepting or rejecting a trade that has already been executed. In the particular embodiment, RAV313 receives trade information from messages generated by persist component 317 and so will not be aware of trades handled by the trading system until they are persisted. In some embodiments, RAV 313 may also be coupled to clearing processes 303 to determine whether trades executed outside the trading system will be accepted/rejected for clearing

[0080] Match engine 315 subscribes to messages related to new orders, order changes, price changes, new triggers, high/low limits and settlements in the example. Match engine publishes information related to executed trades as well as market data related to the trading that goes on within the contract cluster corresponding to the particular match engine instance. Leg pricing component 411 is instantiated within match engine 315 as needed.

[0081] Ticker plant 305 and publish component 405 subscribe to messages containing information that they use to prepare market data. In general, ticker plant 305 may subscribe to messages immediately, before they are persisted whereas publish component 405 subscribes to messages related to persisted events. In this manner, publisher 405, which may provide market data to third party vendors, does not publish any market data related to or derived from orders, trades, or other events that have not yet been made persistent.

[0082] Fig. 5 illustrates the system of Fig. 4 and highlights a series of inter-component message exchanges that occur when an order request is received by validator 311. The bold lines represent an exemplary flow of messages. For example, validator 311 publishes a new order message to which RAV component 313 subscribes. RAV component 313 publishes a new order message to which match engine 315 subscribes. Although it is possible for both match engine 315 and RAV component 313 to subscribe to the same message from validator 311, by imposing a sequential order on the messages match engine 315 will not become aware of a new order until it is accepted as satisfying the business rules implemented by RAV 313. Match engine 315, in turn, publishes a

new order message to persist component 317 such that database 319 holds a persistent copy of each new order.

5 [0083] Referring to Fig. 6, when match engine 315 makes a trade, a trade message is published to subscribers including ticker plant 305 and persist component 317. Persist component 317 stores information about the executed trade in database 319, and updates corresponding order information that was previously stored for the orders related to the executed trade. Persist component 317 publishes a message that is subscribed to by both publish component 405 and RAV component 313. Publish component 405 uses the trade message to compile market data for distribution. RAV component 313 uses the
10 trade message to update account records for use in assessing subsequent orders made on the same account.

[0084] It can be appreciated that the rich and flexible messaging implemented by EMA bus 115 enables the various processes involved in a trading system to be implemented as a plurality of discrete, special-purpose components. Failure of any one component will
15 not cause a failure of other components in the system. This allows the other components to continue operation, or to shut down and restart gracefully so that trade information is not corrupted.

[0085] Although the invention has been described and illustrated with a certain degree of particularity, it is understood that the present disclosure has been made only by way
20 of example, and that numerous changes in the combination and arrangement of parts can be resorted to by those skilled in the art without departing from the spirit and scope of the invention, as hereinafter claimed.